

Exploiting spatial sparsity for multi-wavelength imaging in optical interferometry

Éric Thiébaud¹, Ferréol Soulez¹ and Loïc Denis²

¹ Université de Lyon, Lyon, F-69003, France; Université Lyon 1,
Observatoire de Lyon, 9 avenue Charles André,

Saint-Genis Laval, F-69230, France; CNRS, UMR 5574,
Centre de Recherche Astrophysique de Lyon; École Normale Supérieure de Lyon, Lyon, F-69007, France.

² Université de Lyon, F-42023, Saint-Etienne, France,
CNRS, UMR5516, Laboratoire Hubert Curien, F-42000, Saint-Etienne, France,
Université de Saint-Etienne, Jean Monnet, F-42000, Saint-Etienne, France.

Optical interferometers provide multiple wavelength measurements. In order to fully exploit the spectral and spatial resolution of these instruments, new algorithms for image reconstruction have to be developed. Early attempts to deal with multi-chromatic interferometric data have consisted in recovering a gray image of the object or independent monochromatic images in some spectral bandwidths. The main challenge is now to recover the full 3-D (spatio-spectral) brightness distribution of the astronomical target given all the available data. We describe a new approach to implement multi-wavelength image reconstruction in the case where the observed scene is a collection of point-like sources. We show the gain in image quality (both spatially and spectrally) achieved by globally taking into account all the data instead of dealing with independent spectral slices. This is achieved thanks to a regularization which favors spatial sparsity and spectral grouping of the sources. Since the objective function is not differentiable, we had to develop a specialized optimization algorithm which also accounts for non-negativity of the brightness distribution.

© 2012 Optical Society of America

OCIS codes: 100.3175, 100.3190, 100.4145

1. Context

The objective of stellar interferometric imaging is to recover an approximation of the specific brightness distribution $I_\lambda(\theta)$ of observed astronomical objects given measurements providing incomplete samples of the spatial Fourier transform of $I_\lambda(\theta)$. Reconstruction of a monochromatic image from optical interferometry data is a challenging task which has been the subject of fruitful research and resulted in various algorithms (*e.g.*, Mira [1], BSMEM [2, 3], WISARD [4], the *building-block method* [5]). When dealing with multi-spectral data, a first possibility is to process each wavelength independently and reconstruct a monochromatic image for each subset of measurements from a given spectral channel. For instance, this is what have been done by le Bouquin et al. [6] for the multi-spectral images of the Mira star T Lep. Another possibility is to exploit some assumed spectral continuity of $I_\lambda(\theta)$ and process the multi-spectral data globally to reconstruct an approximation of the 3-D distribution $I_\lambda(\theta)$. This computationally more challenging approach can potentially lead to better reconstructions. Significant improvements have been shown when following such spatio-spectral processing in the context of integral field spectral spectroscopy [7–9]. This paper describes a method to jointly reconstruct multi-spectral optical interferometric data.

In order to simplify the problem, we restricted our

study to the cases where the complex visibilities are observed and where the observed scene is a collection of point-like sources. This correspond, for instance, to the science case of the instrument GRAVITY which will be installed at the Very Large Telescope Interferometer (VLTI) to carry out astrometry with absolute phase reference of stars in the galactic center or in globular clusters [10]. In some sense, this latter assumption makes our algorithm a successor of the CLEAN algorithm [11, 12] and the *building-block method* [5] developed for recovering monochromatic images from radio and optical interferometric data respectively. The CLEAN algorithm have been proposed for the processing of GRAVITY interferometric data [13] but only considering “gray” data and not more than three stars in the field of view. In addition to processing multi-variate data, we also introduce the explicit minimization of a non-differentiable regularization term so as to favor spatial sparsity of the reconstructed brightness distribution in a way which is known to be more efficient [14, 15] than greedy algorithms like CLEAN [16] or the *building-block method* [5]. The method presented in this paper improves over early developments presented as an invited paper at the 2012 SPIE Conf. on Astronomical Telescopes & Instrumentation in Amsterdam [17].

Our paper is organized as follows: we first summarize the inverse approach for image reconstruction from interferometric data and discuss various possibilities to

impose spatial sparsity, we then detail our algorithm for minimizing the objective function; finally we present some results on simulated data and discuss the advantages of our approach.

2. Method

A. General principle of image reconstruction

Following an inverse approach, we state image reconstruction as a constrained optimization problem [18]:

$$\mathbf{x}^+ = \arg \min_{\mathbf{x} \in \mathbb{X}} \{f_{\text{data}}(\mathbf{x}) + \mu f_{\text{prior}}(\mathbf{x})\} \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^{|\mathbf{x}|}$ are the sought image parameters, $|\mathbf{x}| = \text{Card}(\mathbf{x})$ is the number of parameters, $\mathbb{X} \subset \mathbb{R}^{|\mathbf{x}|}$ is the subset of feasible parameters, $f_{\text{data}}(\mathbf{x})$ is a data fitting term which enforces agreement of the model with the measurements $\mathbf{y} \in \mathbb{R}^{|\mathbf{y}|}$, $f_{\text{prior}}(\mathbf{x})$ is a regularization term and $\mu > 0$ is a so-called hyper-parameter used to tune the relative weight of the regularization. Following a Bayesian interpretation, f_{data} is the opposite of the log-likelihood of the measurements, f_{prior} is the opposite of the log of the prior distribution of the parameters, and \mathbf{x}^+ defines the maximum a posteriori (MAP) estimate.

Constraining the solution to belong to the feasible set \mathbb{X} is a mean to impose *strict constraints* such as the non-negativity:

$$\mathbb{X} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq 0\} \quad (2)$$

where the inequality $\mathbf{x} \geq 0$ is to be taken element-wise.

B. Direct model and likelihood

In this study, we assume that the optical interferometric data consist in Fourier transform of the brightness distribution $\hat{I}_\lambda(\boldsymbol{\nu})$ measured for a finite set of spatial frequencies $\boldsymbol{\nu} = \mathbf{B}/\lambda$ with \mathbf{B} the interferometric baseline (projected in a plane perpendicular to the line of sight) and λ the wavelength [19].

The most practical representation of a multi-variate distribution such as $I_\lambda(\boldsymbol{\theta})$ by a finite number of parameters \mathbf{x} consists in sampling $I_\lambda(\boldsymbol{\theta})$ separately along its spatial and spectral dimensions. The image parameters are then:

$$x_{n,\ell} \approx I_{\lambda_\ell}(\boldsymbol{\theta}_n) \quad (3)$$

for $\lambda_\ell \in \mathbb{W}$ the list of sampled wavelengths and $\boldsymbol{\theta}_n \in \mathbb{A}$ the list of angular directions, the so-called *pixels*.

For the sake of notational simplicity, we use the same wavelengths in \mathbb{W} as the ones of the data spectral channels and we denote by $y_{p,m,\ell}$ the real ($p = 1$) or imaginary ($p = 2$) part of the complex visibility obtained with m th baseline in ℓ th spectral channel. This notation is intended to clarify the equations and does not impose or assume that all baselines have been observed in all spectral channels. By considering that complex numbers are just pairs of real values, our notation also avoids dealing

with complex arithmetic. In our framework, the model of the data is affine:

$$\begin{aligned} y_{p,m,\ell} &= (\mathbf{H} \cdot \mathbf{x})_{p,m,\ell} + e_{p,m,\ell} \\ &= \sum_n H_{p,m,n,\ell} x_{n,\ell} + e_{p,m,\ell}, \end{aligned} \quad (4)$$

where the term \mathbf{e} accounts for noise and modeling approximations. Formally, the coefficients of the operator \mathbf{H} are given by [19]:

$$H_{p,m,n,\ell} = \begin{cases} +\cos(\boldsymbol{\theta}_n^\top \cdot \mathbf{B}_m / \lambda_\ell) & \text{for } p = 1 \\ -\sin(\boldsymbol{\theta}_n^\top \cdot \mathbf{B}_m / \lambda_\ell) & \text{for } p = 2 \end{cases} \quad (5)$$

with \mathbf{B}_m the m th observed baseline and $\boldsymbol{\theta}_n^\top \cdot \mathbf{B}_m$ the usual scalar product between \mathbf{B}_m and $\boldsymbol{\theta}_n$.

At least because of the strict constraints imposed by the feasible set \mathbb{X} , solving the image reconstruction problem in Eq. (1) must be carried out by an iterative algorithm. Owing to the size of the problem, a fast version of \mathbf{H} has to be implemented. First, we note that the model is separable along the spectral dimension (using a conventional matrix representation, \mathbf{H} would have a block diagonal structure):

$$\mathbf{y}_\ell = \mathbf{H}_\ell \cdot \mathbf{x}_\ell + \mathbf{e}_\ell, \quad (6)$$

where the index ℓ denotes the sub-vector or the sub-operator restricted to the coefficients corresponding to the ℓ th spectral channel. With the generalization of multi-processor computers or multi-core processors, this property of the operator \mathbf{H} may be easily exploited to parallelize the code to apply \mathbf{H} (or its adjoint \mathbf{H}^\top) to a given argument. Second, an algorithm such as the *nonuniform fast Fourier transform* (NU-FFT) [20] can be implemented to speed up the computations by approximating the operator \mathbf{H}_ℓ by:

$$\mathbf{H}_\ell \approx \mathbf{R}_\ell \cdot \mathbf{F} \cdot \mathbf{S} \quad (7)$$

where \mathbf{F} is the discrete Fourier transform (DFT), \mathbf{R}_ℓ interpolates the discrete spatial frequencies resulting from the DFT at the frequencies observed in ℓ th channel and \mathbf{S} is a zero-padding and apodizing operator. Zero-padding improves the accuracy of the approximation, while apodization pre-compensates for the convolution by the interpolation kernel used in \mathbf{R}_ℓ [20]. Note that only the interpolation in the Fourier domain \mathbf{R}_ℓ depends on the spectral channel. In NU-FFT, \mathbf{S} is diagonal, \mathbf{R}_ℓ is very sparse and \mathbf{F} is implemented by a fast Fourier transform (FFT) algorithm, thus the approximation in Eq. (7) is very fast to compute.

Assuming Gaussian noise distribution, the likelihood term writes:

$$f_{\text{data}}(\mathbf{x}) = \frac{1}{2} (\mathbf{H} \cdot \mathbf{x} - \mathbf{y})^\top \cdot \mathbf{W} \cdot (\mathbf{H} \cdot \mathbf{x} - \mathbf{y}) \quad (8)$$

where $\mathbf{W} \in \mathbb{R}^{|\mathbf{y}| \times |\mathbf{y}|}$ is a statistical weighting matrix; in principle, \mathbf{W} is the inverse of the covariance matrix of the measurements: $\mathbf{W} = \text{Cov}(\mathbf{y})^{-1}$.

C. Regularization based on spatial sparsity

Due to the voids in the spatial frequencies covered by the observations, the constraints provided by the data alone do not suffice to define a unique image. The prior constraints imposed by $f_{\text{prior}}(\mathbf{x})$ are then required to help choosing a unique solution among all the images that are compatible with the measurements.

In this paper, we focus on a particular type of astronomical targets which consist in a number of point-like sources with different spectral energy distributions. This includes the case of multiple stars, globular clusters, or groups of stars as observed in the center of our galaxy. For such objects, the most effective means to regularize the problem is to favor spatially sparse distributions, *i.e.* images with as few sources as possible on a dark background. In this section, we derive expressions of the regularization term $f_{\text{prior}}(\mathbf{x})$ suitable to favor spatially sparse distributions.

1. Fully separable sparsity

It is now well known that using the ℓ_1 norm as the regularization term is an effective mean to impose the sparsity of the solution while approximating the data [14]. This leads to take $f_{\text{prior}}(\mathbf{x}) = f_{\text{sparse}}(\mathbf{x})$ with:

$$f_{\text{sparse}}(\mathbf{x}) = \|\mathbf{x}\|_1 \stackrel{\text{def}}{=} \sum_{k,\ell} |x_{k,\ell}| = \text{sgn}(\mathbf{x})^\top \cdot \mathbf{x}, \quad (9)$$

where $\text{sgn}(\mathbf{x})$ is the sign function applied element-wise to the parameters \mathbf{x} . When the parameters are non-negative, $\text{sgn}(\mathbf{x}) = \mathbf{1}$ with $\mathbf{1} \stackrel{\text{def}}{=} (1, \dots, 1)^\top$.

The regularization term $f_{\text{sparse}}(\mathbf{x})$ in Eq. (9) is completely separable. In our framework where the model is spectrally separable, the global criterion defined in Eq. (1) is therefore separable along the spectral dimension. Provided data from different spectral channels are statistically independent, the image reconstruction can be solved independently for each spectral channel.

If the wavelength samples λ_ℓ of the discrete model \mathbf{x} of $I_\lambda(\theta)$ do not coincide with the effective wavelengths of the data, spectral interpolation of the model is required to match the observed wavelengths. In this case, a certain spectral correlation is intrinsic to the model and the 3-D image reconstruction has to be performed globally even if the regularization does not impose any kind of spectral continuity.

2. Non-separable spatial-only sparsity

Physically, sources emit light at all wavelengths and we expect better image reconstruction if we can favor restored sources having the same position whatever the wavelength. Clearly, this is not achieved by the regularization $f_{\text{sparse}}(\mathbf{x})$ in Eq. (9) which is fully separable. In order to impose some spectral continuity while favoring spatial sparsity, we consider the following regularization instead [7, 15]:

$$f_{\text{joint}}(\mathbf{x}) = \sum_n \left(\sum_\ell x_{n,\ell}^2 \right)^{1/2} \quad (10)$$

with n the spatial index (pixel) and ℓ the spectral channel. The fact that such a regularization favors spatial sparsity and spectral grouping is a consequence of the triangular inequality [15]. This sparsity prior is a special case of several recent generalizations such as group LASSO [21], mixed norms [22] or structured sparsity [23].

3. Explicit spectral continuity and gray object

The penalization defined in Eq. (10) can be seen as a spatial regularization term which favors spectral grouping but no spectral continuity nor spectral smoothness. Some authors [7, 8, 24] have shown the efficiency of exploiting the spectral continuity of the sought distribution by using, in addition to a spatial regularization term, an additional spectral regularization like:

$$f_{\text{spectral}}(\mathbf{x}) = \sum_n \mu_n \sum_\ell (x_{n,\ell+1} - x_{n,\ell})^2 \quad (11)$$

with $\mu_n > 0$ suitable regularization weights. In the limit $\mu_n \rightarrow \infty, \forall n$, the regularization in Eq. (11) amounts to assuming that the spectral energy distributions of all sources are flat, that is:

$$x_{n,\ell} = g_n, \quad \forall (n, \ell) \quad (12)$$

where \mathbf{g} is a *gray image* of the object. To speed up the reconstruction, only the gray image has to be reconstructed, using the model:

$$y_{p,m,\ell} = \sum_n H_{p,m,n,\ell} g_n + e_{p,m,\ell}, \quad (13)$$

and, to impose the spatial sparsity, the ℓ_1 regularization:

$$f_{\text{prior}}(\mathbf{g}) = \|\mathbf{g}\|_1 = \sum_n |g_n| = \text{sgn}(\mathbf{g})^\top \cdot \mathbf{g}. \quad (14)$$

D. Optimization algorithm

Most existing image reconstruction algorithms for optical interferometry (*e.g.*, MiRA [1], BSMEM [2, 3] and WISARD [4]) were designed for minimizing a smooth cost function. For that purpose, non-linear conjugate gradient method [25] or limited memory quasi-Newton methods such as VMLM-B [26] are quite efficient and easy to use as they only require computing the cost function and its gradient. A notable exception is MACIM[27] which is based on a Markov-Chain-Monte-Carlo (MCMC) optimization strategy suitable, in theory, for any type of criteria, in particular the non-smooth and non-convex ones; in practice, this is however too computationally intensive for estimating a large number of parameters as it is the case for image reconstruction. When using non smooth regularizations as the ones in Eq. (9) and Eq. (10) to impose spatial sparsity, optimization algorithms based on Newton method (that is, on a quadratic approximation of the cost function) are inefficient and completely different optimization strategies must be followed to solve the problem in Eq. (1) with non-differentiable cost functions. In our algorithm, we introduce variables splitting [28] to

handle the two terms of the cost function as independently as possible and we implement an *alternating direction method of multipliers* [29] (ADMM) to solve the resulting constrained problem. The augmented Lagrangian with ADMM emerges as the most effective in the family of decomposition methods that includes proximal methods [28], variable splitting with quadratic penalty [30], iterative Bregman [31]. See [32] for detailed comparisons.

1. Variable Splitting and ADMM

Introducing auxiliary variables \mathbf{z} , minimization of the two-term cost function in Eq. (1) can be recast in the equivalent constrained problem:

$$\min_{\mathbf{x} \in \mathbb{X}, \mathbf{z}} \{f_{\text{data}}(\mathbf{z}) + \mu f_{\text{prior}}(\mathbf{x})\} \quad \text{s.t.} \quad \mathbf{x} = \mathbf{z}. \quad (15)$$

Imposing that $\mathbf{x} \in \mathbb{X}$ (*i.e.* $\mathbf{x} \geq \mathbf{0}$) rather than $\mathbf{z} \in \mathbb{X}$ is not arbitrary and our motivation for that choice is explained in what follows.

The augmented Lagrangian [25] is a very useful method to deal with constrained problems such as the one in Eq. (15). In our case, the augmented Lagrangian writes:

$$\mathcal{L}_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = f_{\text{data}}(\mathbf{z}) + \mu f_{\text{prior}}(\mathbf{x}) + \mathbf{u}^\top (\mathbf{x} - \mathbf{z}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z}\|_2^2, \quad (16)$$

with \mathbf{u} the Lagrange multipliers associated with the constraint $\mathbf{x} = \mathbf{z}$, $\rho > 0$ the quadratic weight of the constraints, and $\|\mathbf{v}\|_2$ the Euclidean (ℓ_2) norm of \mathbf{v} . Note that taking $\rho = 0$ yields the classical Lagrangian of the constrained problem.

The alternating direction method of multipliers [29] (ADMM) consists in alternatively minimizing the augmented Lagrangian for \mathbf{x} given \mathbf{z} and \mathbf{u} , then for \mathbf{z} given \mathbf{x} and \mathbf{u} , and finally updating the multipliers \mathbf{u} . This scheme, adapted to our specific problem in Eq. (15), is detailed by the following algorithm with the convention that $\mathbf{v}^{(t)}$ is the value of \mathbf{v} at iteration number t :

Algorithm 1. *Resolution of problem (15) by alternating direction method of multipliers.* Choose initial variables $\mathbf{z}^{(0)}$ and Lagrange multipliers $\mathbf{u}^{(0)}$. Then repeat, for $t = 1, 2, \dots$ until convergence, the following steps:

1. choose $\rho^{(t)} > 0$ and update variables \mathbf{x} :

$$\begin{aligned} \mathbf{x}^{(t)} &= \arg \min_{\mathbf{x} \in \mathbb{X}} \mathcal{L}_{\rho^{(t)}}(\mathbf{x}, \mathbf{z}^{(t-1)}, \mathbf{u}^{(t-1)}) \\ &= \arg \min_{\mathbf{x} \in \mathbb{X}} \left\{ f_{\text{prior}}(\mathbf{x}) + \frac{\rho^{(t)}}{2\mu} \|\mathbf{x} - \tilde{\mathbf{x}}^{(t)}\|_2^2 \right\} \end{aligned} \quad (17)$$

with:

$$\tilde{\mathbf{x}}^{(t)} = \mathbf{z}^{(t-1)} - \mathbf{u}^{(t-1)} / \rho^{(t)}; \quad (18)$$

2. update variables \mathbf{z} :

$$\begin{aligned} \mathbf{z}^{(t)} &= \arg \min_{\mathbf{z}} \mathcal{L}_{\rho^{(t)}}(\mathbf{x}^{(t)}, \mathbf{z}, \mathbf{u}^{(t-1)}) \\ &= \arg \min_{\mathbf{z}} \left\{ f_{\text{data}}(\mathbf{z}) + \frac{\rho^{(t)}}{2} \|\mathbf{z} - \tilde{\mathbf{z}}^{(t)}\|_2^2 \right\} \end{aligned} \quad (19)$$

with:

$$\tilde{\mathbf{z}}^{(t)} = \mathbf{x}^{(t)} + \mathbf{u}^{(t-1)} / \rho^{(t)}; \quad (20)$$

3. update multipliers \mathbf{u} :

$$\mathbf{u}^{(t)} = \mathbf{u}^{(t-1)} + \rho^{(t)} (\mathbf{x}^{(t)} - \mathbf{z}^{(t)}) . \quad \blacksquare \quad (21)$$

Our algorithm can be seen as an instance of SALSAs [33] with however some improvements. First, we deal with the additional constraints that the variables are non-negative. Second, we allow for changing the weight of the augmented penalty at every iteration which can considerably speed-up convergence. Third, for real observations, the operator \mathbf{H} cannot be easily diagonalized (*e.g.* by using FFT) thus the updating of variables \mathbf{z} cannot be exactly carried out. Finally, we consider the possibility of warm starting the algorithm with a solution previously computed. This latter feature is of interest to improve a solution if too few iterations have been performed or to find the solution of the problem with a slightly different value of the hyper-parameter μ .

In Appendix A, we show how the updating of the variables \mathbf{x} (step 1 of Algorithm 1) can be implemented taking into account the constraints that the parameters are non-negative. This is our motivation for imposing $\mathbf{x} \in \mathbb{X}$ on the variables \mathbf{x} and not on the variables \mathbf{z} . This avoids introducing additional auxiliary variables for the sole purpose of accounting for the feasible set. The formulae to update the variables $\mathbf{x} \in \mathbb{X}$ for $f_{\text{sparse}}(\mathbf{x})$ and $f_{\text{joint}}(\mathbf{x})$ are respectively given by Eq. (A5) and Eq. (A13) in Appendix A.

2. Solving for the auxiliary variables

Since $\mathcal{L}_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u})$ is quadratic with respect to \mathbf{z} , updating these variables (step 2 of Algorithm 1) amounts to solving the linear problem:

$$\mathbf{A}^{(t)} \cdot \mathbf{z}^{(t)} = \mathbf{b}^{(t)} \quad (22)$$

with:

$$\mathbf{A}^{(t)} = \mathbf{H}^\top \cdot \mathbf{W} \cdot \mathbf{H} + \rho^{(t)} \mathbf{I} \quad (23)$$

$$\mathbf{b}^{(t)} = \mathbf{H}^\top \cdot \mathbf{W} \cdot \mathbf{y} + \rho^{(t)} \mathbf{x}^{(t)} + \mathbf{u}^{(t-1)} \quad (24)$$

with \mathbf{I} the identity matrix (of suitable size). Since the augmented term is diagonal and provided data from different spectral channels are statistically independent, this problem can be solved separately for each spectral channel:

$$\mathbf{A}_\ell^{(t)} \cdot \mathbf{z}_\ell^{(t)} = \mathbf{b}_\ell^{(t)}, \quad \forall \ell, \quad (25)$$

with:

$$\mathbf{A}_\ell^{(t)} = \mathbf{H}_\ell^\top \cdot \mathbf{W}_\ell \cdot \mathbf{H}_\ell + \rho^{(t)} \mathbf{I} \quad (26)$$

$$\mathbf{b}_\ell^{(t)} = \mathbf{H}_\ell^\top \cdot \mathbf{W}_\ell \cdot \mathbf{y}_\ell + \rho^{(t)} \mathbf{x}_\ell^{(t)} + \mathbf{u}_\ell^{(t-1)} \quad (27)$$

using the same conventions as in Eq. (6).

In practice, we (approximately) solve these problems by means of the conjugate gradients algorithm [25, 34] and starting with the previous solution $\mathbf{z}^{(t-1)}$. According to Eckstein-Bertsekas theorem [35], the ADMM algorithm is proved to converge provided that approximations in the update of auxiliary variables \mathbf{z} be absolutely summable. That is:

$$\sum_{t=1}^{\infty} \|\mathbf{z}^{(t)} - \mathbf{z}_{\text{exact}}^{(t)}\|_2 < \infty \quad (28)$$

must hold with $\mathbf{z}_{\text{exact}}^{(t)}$ the exact solution of Eq. (19) and $\mathbf{z}^{(t)}$ the approximate solution of Eq. (19) returned by the conjugate gradient iterations. In Appendix B, we show how to set the stopping criterion of the conjugate gradient method so that the constraint in Eq. (28) holds. In our tests, we however simply stop the conjugate gradient iterations when the Euclidean norm of the residuals of Eq. (22) becomes significantly smaller than its initial value:

$$\|\mathbf{A}^{(t)} \cdot \mathbf{z}^{(t)} - \mathbf{b}^{(t)}\|_2 \leq \epsilon_{\text{CG}} \|\mathbf{A}^{(t)} \cdot \mathbf{z}^{(t-1)} - \mathbf{b}^{(t)}\|_2 \quad (29)$$

with $\epsilon_{\text{CG}} \in (0, 1)$. For our tests, we took $\epsilon_{\text{CG}} = 10^{-2}$ and allowed a maximum of 5 conjugate gradient iterations. With this simple prescription, we did not experiment any divergence of the global algorithm although it may depend on the problem at hands.

3. Stopping criteria

At the solution $\{\mathbf{x}^*, \mathbf{z}^*, \mathbf{u}^*\}$ of problem (15), Karush-Kuhn-Tucker (KKT) conditions of optimality [25] stipulate that, the constraints must be satisfied and that the solution must be a stationary point of the Lagrangian:

$$\mathbf{x}^* = \mathbf{z}^* \quad (30)$$

$$\mathbf{0} \in \partial_{\mathbf{x}} \mathcal{L}_0(\mathbf{x}^*, \mathbf{z}^*, \mathbf{u}^*) = \mu \partial f_{\text{prior}}(\mathbf{x}^*) + \mathbf{u}^* \quad (31)$$

$$\mathbf{0} \in \partial_{\mathbf{z}} \mathcal{L}_0(\mathbf{x}^*, \mathbf{z}^*, \mathbf{u}^*) = \partial f_{\text{data}}(\mathbf{z}^*) - \mathbf{u}^* \quad (32)$$

where ∂ denotes the subdifferential operator [29]. Since f_{data} is differentiable, ϵ and ∂f_{data} can be replaced by ϵ and by ∇f_{data} , the gradient of f_{data} in the third KKT condition (32) which becomes:

$$\mathbf{u}^* = \nabla f_{\text{data}}(\mathbf{z}^*). \quad (33)$$

If $\mathbf{z}^{(t)}$ exactly minimizes $\mathcal{L}_{\rho^{(t)}}(\mathbf{x}^{(t)}, \mathbf{z}, \mathbf{u}^{(t-1)})$, we have:

$$\begin{aligned} \nabla f_{\text{data}}(\mathbf{z}^{(t)}) - \mathbf{u}^{(t-1)} + \rho^{(t)}(\mathbf{z}^{(t)} - \mathbf{x}^{(t)}) &= \mathbf{0} \\ \implies \nabla f_{\text{data}}(\mathbf{z}^{(t)}) &= \mathbf{u}^{(t)}, \end{aligned} \quad (34)$$

thus the 3rd KKT condition in Eq. (33) is automatically satisfied at the end of an *exact* ADMM iteration.

As we use the method of conjugate gradients to solve Eq. (22), Eq. (34) is only approximately satisfied. Moreover, updating the multipliers \mathbf{u} according to step 3 of the algorithm may be subject to accumulation of rounding errors. The stability of the algorithm or its convergence rate may be improved by taking $\mathbf{u}^{(t)} = \nabla f_{\text{data}}(\mathbf{z}^{(t)})$. In

our tests, though, we have not seen significant differences between updating the Lagrange multipliers according to Eq. (21) or according to Eq. (34).

Since $\mathbf{x}^{(t)}$ minimizes $\mathcal{L}_{\rho^{(t)}}(\mathbf{x}, \mathbf{z}^{(t-1)}, \mathbf{u}^{(t-1)})$, we have:

$$\begin{aligned} \mathbf{0} &\in \mu \partial f_{\text{prior}}(\mathbf{x}^{(t)}) + \mathbf{u}^{(t-1)} + \rho^{(t)}(\mathbf{x}^{(t)} - \mathbf{z}^{(t-1)}) \\ &\in \mu \partial f_{\text{prior}}(\mathbf{x}^{(t)}) + \mathbf{u}^{(t)} + \rho^{(t)}(\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}) \\ \implies -\rho^{(t)}(\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}) &\in \mu \partial f_{\text{prior}}(\mathbf{x}^{(t)}) + \mathbf{u}^{(t)} \end{aligned}$$

thus:

$$\mathbf{s}^{(t)} = \rho^{(t)}(\mathbf{z}^{(t)} - \mathbf{z}^{(t-1)}) \quad (35)$$

can be seen as the residuals for the 2nd KKT condition in Eq. (31), while:

$$\mathbf{r}^{(t)} = \mathbf{x}^{(t)} - \mathbf{z}^{(t)} \quad (36)$$

are the residuals for the primary constraint in Eq. (30).

Finally, the KKT conditions imply that the so-called *primal* and *dual residuals* [29] defined in Eq. (36) and Eq. (35) must converge to zero. Following [29], we therefore stop the algorithm when:

$$\|\mathbf{r}^{(t)}\|_2 \leq \tau_{\text{prim}}^{(t)} \quad \text{and} \quad \|\mathbf{s}^{(t)}\|_2 \leq \tau_{\text{dual}}^{(t)}, \quad (37)$$

where the convergence thresholds are given by:

$$\tau_{\text{prim}}^{(t)} \stackrel{\text{def}}{=} \sqrt{N} \epsilon_{\text{abs}} + \epsilon_{\text{rel}} \max(\|\mathbf{x}^{(t)}\|_2, \|\mathbf{z}^{(t)}\|_2), \quad (38)$$

$$\tau_{\text{dual}}^{(t)} \stackrel{\text{def}}{=} \sqrt{N} \epsilon_{\text{abs}} + \epsilon_{\text{rel}} \|\mathbf{u}^{(t)}\|_2, \quad (39)$$

where $N = \text{Card}(\mathbf{x})$ is the number of sought parameters, $\epsilon_{\text{abs}} \geq 0$ and $\epsilon_{\text{rel}} \in (0, 1)$ are absolute and relative convergence tolerances. For our tests, we found that $\epsilon_{\text{abs}} = 0$ and $\epsilon_{\text{rel}} = 10^{-3}$ yield sufficient precision for the solution.

4. Initialization and warm start

Given an initial estimate $\mathbf{z}^{(0)}$ for the auxiliary variables, the 3rd KKT condition in Eq. (33) suggests to start the iterative algorithm with initial Lagrange multipliers given by:

$$\mathbf{u}^{(0)} = \nabla f_{\text{data}}(\mathbf{z}^{(0)}) = \mathbf{H}^\top \cdot \mathbf{W} \cdot (\mathbf{H}^\top \cdot \mathbf{z}^{(0)} - \mathbf{y}). \quad (40)$$

Starting the algorithm as suggested has several advantages. First, there are no needs for the initial variables $\mathbf{z}^{(0)}$ to belong to the feasible set. Second, this readily provides initial Lagrange multipliers. Note that starting with an initial estimate $\mathbf{x}^{(0)}$ for the variables \mathbf{x} would have the double drawback that $\mathbf{x}^{(0)}$ must be feasible and that since $f_{\text{prior}}(\mathbf{x})$ may be non-differentiable it does not yields an explicit expression for the initial Lagrange multipliers.

The other required initial setting is the value of the augmented penalty parameter $\rho^{(1)}$ used to compute the first estimate $\mathbf{x}^{(1)}$ of the variables \mathbf{x} given $\mathbf{z}^{(0)}$ and $\mathbf{u}^{(0)}$. For the subsequent iterations, ρ can be kept constant or updated according to the prescription described in Sec. 2 D 5.

To continue the iterations or compute a solution with slightly different parameters (*e.g.* the regularization parameter μ), the possibility to restart the algorithm with the output of a previous run with no loss of performances regarding the rate of convergence is a needed feature. This is called *warm restart* and is simply achieved by saving a minimal set of variables upon return of the algorithm. Since each iteration of our algorithm starts by computing the variables \mathbf{x} given the auxiliary variables \mathbf{z} , the Lagrange multipliers \mathbf{u} and the augmented penalty parameter ρ , it is sufficient to save $\{\mathbf{z}, \mathbf{u}, \rho\}$ for being able to warm restart the method.

5. Tuning the augmented penalty parameter ρ

One of the important settings of the ADMM method is the value of the augmented penalty parameter ρ : if it is too small, the primal constraints $\mathbf{x} = \mathbf{z}$ will converge too slowly; while the cost functions will decrease too slowly if ρ is too large. Some authors, *e.g.* [33], use a constant augmented penalty parameter for all the iterations which requires trials and errors to find an efficient value for ρ . In fact, it is worth using a good value of ρ at every iteration of ADMM to accelerate the convergence [29]. In this section, we describe means to automatically derive a suitable value for the augmented penalty parameter following a simple reasoning.

The convergence criterion defined in Eq. (37) is equivalent to have:

$$\phi^{(t)} \leq 1 \quad \text{with} \quad \phi^{(t)} \stackrel{\text{def}}{=} \max \left(\frac{\|\mathbf{r}^{(t)}\|_2}{\tau_{\text{prim}}^{(t)}}, \frac{\|\mathbf{s}^{(t)}\|_2}{\tau_{\text{dual}}^{(t)}} \right). \quad (41)$$

According to the updating rules in one ADMM iteration (see Algorithm 1), $\phi^{(t)}$ does only depend on $\mathbf{z}^{(t-1)}$, $\mathbf{u}^{(t-1)}$ and $\rho^{(t)}$. The augmented penalty parameter $\rho^{(t)}$ is therefore the only tunable parameter that has an incidence on the value of $\phi^{(t)}$ for the t^{th} ADMM iteration. The idea is then to chose the value of $\rho^{(t)}$ so as to approximately minimize $\phi^{(t)}$. In terms of number of ADMM iterations, we expect to achieve the faster convergence of the algorithm in that way. However, tuning ρ at every ADMM iteration requires to repeat each iteration for different values of ρ and has therefore the same computational cost as several ADMM iterations. A compromise has to be found between the accuracy on ρ and the number of trials.

Our objective is to derive an economical way to find:

$$\rho^{(t)} \approx \rho_*^{(t)} \stackrel{\text{def}}{=} \arg \min_{\rho} \phi_t(\rho). \quad (42)$$

where $\phi_t(\rho)$ is the value taken by $\phi^{(t)}$ when $\rho^{(t)} = \rho$. All the quantities, $\|\mathbf{r}^{(t)}\|_2$, $\|\mathbf{s}^{(t)}\|_2$, $\tau_{\text{prim}}^{(t)}$, and $\tau_{\text{dual}}^{(t)}$, involved in $\phi^{(t)}$ vary continuously (though not necessarily smoothly) with respect to $\rho^{(t)}$; hence, considering the definition of $\rho_*^{(t)}$ and $\phi^{(t)}$, we obtain the following implication:

$$\rho^{(t)} = \rho_*^{(t)} \implies \frac{\|\mathbf{r}^{(t)}\|_2}{\tau_{\text{prim}}^{(t)}} = \frac{\|\mathbf{s}^{(t)}\|_2}{\tau_{\text{dual}}^{(t)}}. \quad (43)$$

Besides, the norm of the primal residuals $\|\mathbf{r}^{(t)}\|_2$ is a decreasing function of $\rho^{(t)}$; while the norm of the dual residuals $\|\mathbf{s}^{(t)}\|_2$ is an increasing function of $\rho^{(t)}$ [29] and, close enough to the solution, the values of τ_{prim} and τ_{dual} should converge to their final values and thus not depend too much on ρ . Under these assumptions, the ratios $\|\mathbf{r}^{(t)}\|_2/\tau_{\text{prim}}^{(t)}$ and $\|\mathbf{s}^{(t)}\|_2/\tau_{\text{dual}}^{(t)}$ should also be decreasing and increasing functions of ρ respectively. Close to the solution $\{\mathbf{x}^*, \mathbf{z}^*, \mathbf{u}^*\}$ of the problem, the necessary condition in Eq. (43) is therefore also a sufficient condition to define the optimal value $\rho_*^{(t)}$. These considerations lead us to choose $\rho^{(t)}$ such that:

$$\eta^{(t)} \approx 1 \quad \text{with} \quad \eta^{(t)} \stackrel{\text{def}}{=} \frac{\|\mathbf{r}^{(t)}\|_2 \tau_{\text{dual}}^{(t)}}{\|\mathbf{s}^{(t)}\|_2 \tau_{\text{prim}}^{(t)}}, \quad (44)$$

which is expected to be a decreasing function of $\rho^{(t)}$ close to the solution. A better alternative may be to choose $\rho^{(t)}$ such that:

$$\eta_{\text{alt}}^{(t)} \approx 1 \quad \text{with} \quad \eta_{\text{alt}}^{(t)} \stackrel{\text{def}}{=} \frac{\|\mathbf{r}^{(t)}\|_2 \tau_{\text{dual}}^{(t-1)}}{\|\mathbf{s}^{(t)}\|_2 \tau_{\text{prim}}^{(t-1)}}. \quad (45)$$

Indeed, as $\tau_{\text{dual}}^{(t-1)}$ and $\tau_{\text{prim}}^{(t-1)}$ do not depend on $\rho^{(t)}$, $\eta_{\text{alt}}^{(t)}$ is always decreasing function of $\rho^{(t)}$ while approaching $\eta^{(t)}$ when the algorithm is close to the solution. The following algorithm implements our safeguarded strategy to find $\rho^{(t)} > 0$ such that $\eta^{(t)} \approx 1$ or $\eta_{\text{alt}}^{(t)} \approx 1$.

Algorithm 2. *Tuning of the augmented penalty parameter ρ so that $\eta \approx 1$. Choose $\sigma \geq 0$, $\tau > 1$, $\gamma > 1$ and an initial value for ρ , and set $\rho_{\min} = 0$, $\rho_{\max} = +\infty$. Then, until convergence, repeat the following steps:*

1. Update \mathbf{x} , \mathbf{z} , and \mathbf{u} according to ADMM updating rules. Compute η , defined in Eq. (44) or in Eq. (45), and ϕ , defined in Eq. (41).
2. If $1/\tau \leq \eta \leq \tau$ or $\phi < \sigma \phi^{(t-1)}$, accept the current solution and stop.
3. If $\eta < 1/\tau$, then ρ is too large; let $\rho_{\max} := \rho$ and

$$\rho := \begin{cases} \sqrt{\rho_{\min} \rho_{\max}} & \text{if } \rho_{\min} > 0 \\ \rho_{\max}/\gamma & \text{otherwise} \end{cases}$$

then go to step 1.

4. If $\eta > \tau$, then ρ is too small; let $\rho_{\min} := \rho$ and

$$\rho := \begin{cases} \sqrt{\rho_{\min} \rho_{\max}} & \text{if } \rho_{\max} < \infty \\ \gamma \rho_{\min} & \text{otherwise} \end{cases}$$

then go to step 1. ■

The following remarks clarify some aspects of this algorithm:

- To simplify the notations, we dropped the index t of the ADMM iteration in the equations of Algorithm 2. The updating of variables (step 1) must be understood as computing $\mathbf{x}^{(t)}$, $\mathbf{z}^{(t)}$, *etc.* given $\mathbf{z}^{(t-1)}$ and $\mathbf{u}^{(t-1)}$, and assuming $\rho^{(t)} = \rho$.
- Except for the very first ADMM iteration ($t = 1$), the initial value for ρ is the previous selected value $\rho^{(t-1)}$. For the first iteration, we derive an initial value of ρ such that:

$$\begin{aligned} \rho^{(1)} &= \arg \min_{\rho} f_{\text{data}}(\mathbf{z}^{(1)} - \mathbf{u}^{(1)}/\rho) \\ &= \frac{\mathbf{u}^{(1)\top} \cdot \mathbf{H}^\top \cdot \mathbf{W} \cdot \mathbf{H} \cdot \mathbf{u}^{(1)}}{\mathbf{u}^{(1)\top} \cdot \mathbf{u}^{(1)}} \end{aligned} \quad (46)$$

with $\mathbf{u}^{(1)} = \nabla f_{\text{data}}(\mathbf{z}^{(1)}) = \mathbf{H}^\top \cdot \mathbf{W} \cdot (\mathbf{H} \cdot \mathbf{z}^{(1)} - \mathbf{y})$ and which amounts to have the best $\tilde{\mathbf{x}}^{(1)}$, defined in Eq. (18), with respect to f_{data} . This choice has the advantage of avoiding an initialization with an arbitrary value for ρ . The rule does however not yield an efficient strategy for tuning ρ at every iteration.

- Algorithm 2 is safeguarded in the sense that it maintains a strict bracketing $\rho_{\min} < \rho_*^{(t)} < \rho_{\max}$ of the solution.
- In step 2 of Algorithm 2: The value of ρ is accepted when $1/\tau \leq \eta \leq \tau$ which, with $\tau > 1$, is how we express that $\eta \approx 1$. We achieved good results with $\tau = 1.2$ in our tests. The current value of ρ is also accepted, if the relative reduction in the convergence criterion $\phi^{(t)}$, defined in Eq. (41), is better than σ with respect to the previous iteration. This shortcut helps to reduce the number of inner iterations. To avoid this shortcut, it is sufficient to take: $\sigma = 0$. We took $\sigma = 0.9$ in our tests.
- In step 3 and step 4 of Algorithm 2: When ρ has been bracketed by $(\rho_{\min}, \rho_{\max})$, taking $\rho = \sqrt{\rho_{\min} \rho_{\max}}$, that is the geometrical means of the end points, is similar to a bisection step in a zero finding algorithm.
- For the first ADMM iteration, the magnitude of ρ is not yet known so to avoid too many iterations, we use a larger value of the loop gain γ , say $\gamma = 10$ when $t = 1$ and $\gamma = 1.5$ for $t > 1$.

Another possibility is to always accept an ADMM iteration and simply use the value of $\eta^{(t)}$ to determine whether ρ should be reduced, kept the same, or augmented for the next iteration. For instance:

$$\rho^{(t+1)} = \begin{cases} \gamma \rho^{(t)} & \text{if } \eta_t > \tau, \\ \rho^{(t)}/\gamma & \text{if } \eta_t < 1/\tau, \\ \rho^{(t)} & \text{else;} \end{cases} \quad (47)$$

with $\tau \geq 1$ and $\gamma > 1$. In words, ρ is augmented (by multiplying it by a factor γ) whenever the relative size

of the primal residuals is significantly larger than that of the dual residuals; while ρ is reduced (by dividing it by a factor γ) whenever the relative size of the primal residuals is significantly smaller than that of the dual residuals. This strategy is similar to the one described by Boyd *et al.* [29] except that our prescription properly scales with the magnitudes of the residuals and of the objective function of the problem so we expect a better behavior.

E. Debiasing the solution

One of the drawback of sparsity imposed by means of the ℓ_1 norm is that it yields a result which is biased toward zero [36]. This is shown by the recovered spectra in Fig. 2 and by the brightness distributions of the pixels in Fig. 3. Since the sparsity constraint really improves the detection of the sources, the resulting image can be used to decide where the sources are. By thresholding the gray image or the wavelength integrated multi-spectral image resulting from a reconstruction with a spatial sparsity constraint, we define a sparse spatial support \mathcal{S} containing all detected sources. Then, as proposed in [37], for the debiasing step, we minimize the likelihood function $f_{\text{data}}(\mathbf{x})$ with non-negative constraints only over $\mathbf{x}_{\mathcal{S}}$ defined as the parameters \mathbf{x} restricted to the support \mathcal{S} while keeping all other parameters equal to zero. As the sub-matrix $\mathbf{H}_{\mathcal{S}}$ containing the columns of \mathbf{H} restricted by \mathcal{S} is well conditioned, no additional prior is needed to define the debiased solution.

3. Results

To check the proposed algorithm, we simulated a cluster of 50 stars with random positions and luminosities and with spectra randomly taken from the library compiled by Jacoby *et al.* [38]. The field of view is 128×128 pixels with 0.5 milliarcseconds/pixel and we took 100 spectral channels from $\lambda = 493 \text{ nm}$ to $\lambda = 507 \text{ nm}$ by steps of $\Delta\lambda = 0.14 \text{ nm}$. To simulate the observations, we took 100 random interferometric baselines with a maximum baseline of 180 m. We added Gaussian white noise to the complex visibilities with a level such that the maximum signal-to-noise ratio (SNR) is equal to 100.

For the image reconstructions, we considered three different cases: the reconstruction of a multi-spectral distribution with the regularization $f_{\text{sparse}}(\mathbf{x})$ in Eq. (9), or the regularization $f_{\text{joint}}(\mathbf{x})$ in Eq. (10), and the reconstruction of a gray object \mathbf{g} with the regularization $f_{\text{sparse}}(\mathbf{g})$. In order to set the relative weight of the priors, we choose the value of the hyper-parameter μ which yields an image which has the least mean square error with the true distribution. Once the values of μ and ρ are chosen, the reconstruction of a $128 \times 128 \times 100$ distribution from $\sim 2 \times 10^4$ measurements takes about 4 minutes on a GNU/Linux workstation with a quad-core processor at 3 GHz and using a multi-threaded version of FFTW [39] to compute the discrete Fourier transforms.

Figure 1 shows the integrated flux, *i.e.* $\sum_{\ell} x_{n,\ell}$, for




Fig. 1. Integrated flux for the star cluster. From top to bottom, left to right: true object; reconstruction with fully separable sparsity prior; reconstruction with spatial sparsity and assuming a gray object; reconstruction with joint-sparsity prior. The spectra of the sources encircled by the boxes are shown in Fig. 2. Axes units are in milliarcseconds.

the true distribution and for the reconstructed ones. In all cases, sparsity priors effectively yield a solution with point-like structures. However, when there is no trans-spectral constraints, only a few sources are correctly found and there are many more spurious sources. When using $f_{\text{joint}}(\mathbf{x})$ or assuming a gray object, the estimated integrated luminosity is much more consistent with that of the true object: all existing sources are found and the spurious sources are not only less numerous but also much fainter than the true ones. This is shown by the brightness distributions depicted by Fig. 3 and Fig. 4.

Figure 2 shows the spectra of the two stars encircled by boxes in Fig. 1, clearly the spectra recovered with $f_{\text{joint}}(\mathbf{x})$ (thin curves marked with open triangles) are of much higher quality than the spectra estimated when treating the spectral channels independently (thin curves marked with open boxes). Compared to the true spectra (thick lines) there is however a small but significant bias in the spectra obtained with $f_{\text{joint}}(\mathbf{x})$. This is not unexpected as the mixed norm implemented by $f_{\text{joint}}(\mathbf{x})$ results in an attenuation as shown by Eq. (A11).

As shown by Fig. 4, in the reconstructed gray image or in the image reconstructed with $f_{\text{joint}}(\mathbf{x})$, all sources whose mean flux is greater than 1 are true positive detections while all false positives have a smaller mean flux. We therefore select the sources with mean fluxes greater than this level to apply the debiasing method described in Sec. 2E to effectively remove this bias as shown by the




Fig. 2. Spectra of two selected sources. Each panel show the spectra of one of the sources encircled by the boxes in Fig. 1. Thick lines are for the true spectra and thin lines with markers for the restored spectra. The open squares indicate the reconstruction with fully separable sparsity prior $f_{\text{sparse}}(\mathbf{x})$; the open triangles indicate the reconstruction with joint sparsity prior $f_{\text{joint}}(\mathbf{x})$; the filled triangles indicate the restored spectra after debiasing (which are virtually indistinguishable from the true ones).

thin curves with filled triangles in Fig. 2.

In our reconstructions with the regularization $f_{\text{joint}}(\mathbf{x})$, we found that $\phi^{(t)} < 10^{-3}$ was a good threshold for the global convergence of the algorithm and we compared the different strategies proposed in Sec. 2D5 to set the augmented penalty parameter ρ . The evolution of the convergence criterion for some of these strategies is plotted in Figure 5. With a constant value for ρ , we observed that the rate of convergence is quite sensitive to the value of the augmented penalty parameter. Indeed with $\rho = 3 \times 10^2$ (which is the best value we found), the algorithm converged in 455 s, while it took 1 071 s and 847 s with $\rho = 10^2$ and $\rho = 10^3$ respectively. Although we did not try many different values for the parameters τ and γ , we found that the automatic strategies for setting ρ with




Fig. 3. Histograms of the mean fluxes of the sources for the true object (in *black*), for the 3-D images restored with fully separable sparsity (in *white*) and with joint-sparsity (in *dark gray*) priors, and for the 2-D gray image restored with sparsity prior (in *light gray*). The vertical scale has been truncated to focus on the distributions of the brightest sources.




Fig. 4. Histograms of the mean flux of the true and false positive detection in the reconstructions under joint sparsity and gray sparsity priors. A positive detection is defined as a pixel with non-zero mean flux in the reconstruction. The vertical scale has been truncated to focus on the distributions of the true positive sources.

Algorithm 2 and $\eta^{(t)}$ or according to Eq. (47) failed with their convergence criterion oscillating with $\phi^{(t)} \approx 10^{-2}$. In fact, in spite of the loss of time due to the number of retries needed to find a correct value for ρ at each ADMM iteration, we found that the best strategy was to use $\eta_{\text{alt}}^{(t)}$ in Algorithm 2. In this case, it seemed to be better to use a tighter tolerance for $\eta_{\text{alt}}^{(t)} \approx 1$ as with $\tau = 3$ the algorithm converged in 402s, while it took only 250s with $\tau = 1.2$ (see Fig. 5).




Fig. 5. Evolution of the convergence criterion $\phi^{(t)}$, defined in Eq. (41), for different strategies to choose the augmented penalty parameter. Dashed curves are for a constant ρ . Solid curves are for ρ automatically set to have $\eta_{\text{alt}}^{(t)} \approx 1$.

4. Discussion and Perspectives

We have shown the importance of using trans-spectral constraints to improve the quality of the restoration of the multi-spectral brightness distribution $I_\lambda(\theta)$ of an astronomical target from optical interferometric data. These results confirm what has been observed for other types of data (like integral field spectroscopy).

For the moment, our demonstration is restricted to specific objects which are spatially sparse (*e.g.* point-like sources) and must be generalized to other types of spatial distributions. Being implemented by non-differentiable cost functions, spatial sparsity requires specific optimization algorithms. We have shown that variable splitting by the alternate direction method of multipliers (ADMM) is suitable to solve the optimization problem in a short amount of time. In addition to being able to deal with non-differentiable criteria, the ADMM method leads to splitting the full problem in sub-problems that are easier to solve and that may be independent. This straightforwardly gives the opportunity of speeding up the code, *e.g.* by means of parallelization. This possibility remains if other priors are used, *e.g.* to account for a smooth spatial distribution.

To simplify the problem at hand, we considered that complex visibilities have been measured. At optical wavelengths, this is only possible with phase referencing [40]. In order to process most existing interferometric data, we will have to modify the likelihood term $f_{\text{data}}(\mathbf{x})$ and use a non-linear method (*i.e.* not the linear conjugate gradients) to update the auxiliary variables \mathbf{z} . The new algorithm that we proposed, because it splits the two cost functions, $f_{\text{data}}(\mathbf{z})$ and $f_{\text{prior}}(\mathbf{x})$, may however be an efficient alternative to the variable metric method used in MiRA [1, 26] or the non-linear conjugate gradients

method in BSMEM [2, 3] which considers directly the sum of the cost functions.

Acknowledgments

This work is supported by the French ANR (*Agence Nationale de la Recherche*), Éric Thiébaud and Loïc Denis work for the MiTiV project (*Méthodes Inverses de Traitement en Imagerie du Vivant*, ANR-09-EMER-008) and Ferréol Soulez is funded by the POLCA project (*Percées astrophysiques grâce au traitement de données interférométriques polychromatiques*, ANR-10-BLAN-0511).

Our algorithm has been implemented and tested with YORICK (<http://yorick.sourceforge.net/>) which is freely available.

Appendix A: Proximity Operators for Spatial Sparsity of Non-Negative Variables

Updating of the variables \mathbf{x} by Eq. (17) and (18) in the ADMM method consists in solving a problem of the form:

$$\min_{\mathbf{x} \in \mathbb{X}} \left\{ \alpha f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2 \right\}, \quad (\text{A1})$$

with $\alpha = \mu/\rho^{(t)} > 0$ and $f(\mathbf{x}) = f_{\text{prior}}(\mathbf{x})$. Solving problem (A1) is very close to applying the so-called proximity operator (also known as *Moreau proximal mapping*) of the function $\alpha f(\mathbf{x})$ which is defined by [28]:

$$\text{prox}_{\alpha f}(\tilde{\mathbf{x}}) \stackrel{\text{def}}{=} \arg \min_{\mathbf{x} \in \mathbb{R}^N} \left\{ \alpha f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2 \right\}. \quad (\text{A2})$$

Proximity operators for non differentiable cost functions like $f_{\text{sparse}}(\mathbf{x})$ or $f_{\text{joint}}(\mathbf{x})$ have already been derived [28] and we simply need to modify them to account for the additional constraint that $\mathbf{x} \in \mathbb{X}$. Since \mathbb{X} is the subset of non-negative vectors of \mathbb{R}^N , *i.e.* $\mathbb{X} = \mathbb{R}_+^N$, we denote by:

$$\text{prox}_{\alpha f}^+(\tilde{\mathbf{x}}) \stackrel{\text{def}}{=} \arg \min_{\mathbf{x} \in \mathbb{R}_+^N} \left\{ \alpha f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2 \right\}, \quad (\text{A3})$$

the modified proximity operator that we use to update the variables \mathbf{x} in our algorithm while accounting for non-negativity.

1. Separable Sparsity

The proximity operator for $f_{\text{sparse}}(\mathbf{x})$ defined in Eq. (9), that is the ℓ_1 -norm of \mathbf{x} , is the so called *soft thresholding* operator [28]:

$$\text{prox}_{\alpha f_{\text{sparse}}}(\tilde{\mathbf{x}})_{n,\ell} = \begin{cases} \tilde{x}_{n,\ell} - \alpha & \text{if } \tilde{x}_{n,\ell} > \alpha; \\ \tilde{x}_{n,\ell} + \alpha & \text{if } \tilde{x}_{n,\ell} < -\alpha; \\ 0 & \text{else.} \end{cases} \quad (\text{A4})$$

Imposing the non-negativity is straightforward and yields:

$$\text{prox}_{\alpha f_{\text{sparse}}}^+(\tilde{\mathbf{x}})_{n,\ell} = \begin{cases} \tilde{x}_{n,\ell} - \alpha & \text{if } \tilde{x}_{n,\ell} > \alpha; \\ 0 & \text{else.} \end{cases} \quad (\text{A5})$$

This shows that if $\alpha \geq \max_{n,\ell} \tilde{x}_{n,\ell}$, the output of the proximity operator is zero everywhere.

2. Spatio-Spectral Regularization

Using the joint-sparsity regularization $f_{\text{joint}}(\mathbf{x})$ given by Eq. (10), we aim at minimizing the criterion:

$$c(\mathbf{x}) = \alpha \underbrace{\sum_n \left(\sum_\ell x_{n,\ell}^2 \right)^{1/2}}_{f_{\text{joint}}(\mathbf{x})} + \frac{1}{2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2$$

which is strictly convex with respect to the variables \mathbf{x} [15]. We note that $c(\mathbf{x})$ is separable with respect to the pixel index n . Thus all computations can be done independently for the spectral energy distribution of each pixel.

Considering first the unconstrained case and for variables \mathbf{x} such that the function $f_{\text{joint}}(\mathbf{x})$ is differentiable, minimizing $c(\mathbf{x})$ with respect to the variables \mathbf{x} amounts to finding the root of the partial derivatives of $c(\mathbf{x})$:

$$\begin{aligned} \frac{\partial c(\mathbf{x})}{\partial x_{n,\ell}} = 0 & \iff \frac{\alpha}{\beta_n} x_{n,\ell} + (x_{n,\ell} - \tilde{x}_{n,\ell}) = 0 \\ & \iff x_{n,\ell} = \frac{\tilde{x}_{n,\ell}}{1 + \alpha/\beta_n}, \end{aligned} \quad (\text{A6})$$

with:

$$\beta_n \stackrel{\text{def}}{=} \left(\sum_\ell x_{n,\ell}^2 \right)^{1/2}, \quad (\text{A7})$$

the Euclidean norm of the spectral energy distribution of the n th pixel of \mathbf{x} . Assuming for the moment that $\beta_n > 0$, otherwise $c(\mathbf{x})$ is not differentiable, combining Eq. (A6) and Eq. (A7) yields:

$$\beta_n = \frac{\tilde{\beta}_n}{1 + \alpha/\beta_n}, \quad (\text{A8})$$

since $\alpha/\beta_n > 0$ and with:

$$\tilde{\beta}_n \stackrel{\text{def}}{=} \left(\sum_\ell \tilde{x}_{n,\ell}^2 \right)^{1/2}. \quad (\text{A9})$$

Solving Eq. (A8) for β_n yields:

$$\beta_n = \tilde{\beta}_n - \alpha. \quad (\text{A10})$$

The non-differentiable case occurs when the above expression yields a value of β_n which is not strictly positive, that is when $\tilde{\beta}_n \leq \alpha$, in which case the minimum

of the cost function is given by $x_{n,\ell} = 0, \forall \ell$. Finally, the proximity operator of $\alpha f_{\text{joint}}(\mathbf{x})$ is:

$$\text{prox}_{\alpha f_{\text{joint}}}(\tilde{\mathbf{x}})_{n,\ell} = \begin{cases} \left(1 - \frac{\alpha}{\tilde{\beta}_n}\right) \tilde{x}_{n,\ell} & \text{if } \tilde{\beta}_n > \alpha; \\ 0 & \text{else} \end{cases} \quad (\text{A11})$$

where $\tilde{\beta}_n$ is the Euclidean norm of the spectral energy distribution of the n th pixel of $\tilde{\mathbf{x}}$ defined by Eq. (A9).

In the differentiable case, requiring that $x_{n,\ell} \geq 0$ yields a simple modification of the unconstrained solution given by Eq. (A6):

$$x_{n,\ell} = \frac{\max(0, \tilde{x}_{n,\ell})}{1 + \alpha/\tilde{\beta}_n}, \quad (\text{A12})$$

since $1 + \alpha/\tilde{\beta}_n > 0$. The rest of the reasoning is similar than the unconstrained case except that $\tilde{\beta}_n$ has to be replaced by $\tilde{\beta}_n^+$ the Euclidean norm of the spectral energy distribution of the n th pixel of $\max(0, \tilde{\mathbf{x}})$. The proximity operator of $f_{\text{joint}}(\mathbf{x})$ modified to account for non-negativity is finally:

$$\text{prox}_{\alpha f_{\text{joint}}}^+(\tilde{\mathbf{x}}) = \text{prox}_{\alpha f_{\text{joint}}}(\max(0, \tilde{\mathbf{x}})). \quad (\text{A13})$$

Appendix B: Stopping criterion for the conjugate gradient method

We derive here a possible strategy to set the stopping criterion for the conjugate gradient method used to update the auxiliary variables \mathbf{z} so as to guarantee the global convergence of the ADMM method. If $\mathbf{z}_{\text{exact}}^{(t)}$ is the solution of the linear system $\mathbf{A}^{(t)} \cdot \mathbf{z} = \mathbf{b}^{(t)}$, then:

$$\mathbf{A}^{(t)} \cdot \mathbf{z}_{\text{exact}}^{(t)} - \mathbf{b}^{(t)} = \mathbf{0};$$

while, for the approximate solution:

$$\mathbf{A}^{(t)} \cdot \mathbf{z}^{(t)} - \mathbf{b}^{(t)} = \mathbf{v}^{(t)},$$

where $\mathbf{v}^{(t)}$ are the so called residuals at the end of the conjugate gradients iterations. Then:

$$\mathbf{z}^{(t)} - \mathbf{z}_{\text{exact}}^{(t)} = [\mathbf{A}^{(t)}]^{-1} \cdot \mathbf{v}^{(t)}.$$

As $\mathbf{A}^{(t)} = \mathbf{H}^\top \cdot \mathbf{W} \cdot \mathbf{H} + \rho^{(t)} \mathbf{I}$ with $\rho^{(t)} > 0$ and since $\mathbf{H}^\top \cdot \mathbf{W} \cdot \mathbf{H}$ is at least positive semi-definite, the smallest eigenvalue of $\mathbf{A}^{(t)}$ is greater or equal $\rho^{(t)}$, thus:

$$\|\mathbf{z}^{(t)} - \mathbf{z}_{\text{exact}}^{(t)}\|_2 \leq \|\mathbf{v}^{(t)}\|_2 / \rho^{(t)}.$$

To have:

$$\sum_{t=1}^{\infty} \|\mathbf{z}^{(t)} - \mathbf{z}_{\text{exact}}^{(t)}\|_2 < \infty$$

a sufficient condition is therefore to make sure that:

$$\sum_{t=1}^{\infty} \|\mathbf{v}^{(t)}\|_2 / \rho^{(t)} < \infty.$$

This can be achieved by imposing at each iteration that the stopping criterion for the conjugate gradients be such that:

$$\|\mathbf{v}^{(t)}\|_2 \leq \gamma_{\text{CG}} \rho^{(t)} \xi_{\text{CG}}^t \quad (\text{B1})$$

with $\gamma_{\text{CG}} > 0$ and $\xi_{\text{CG}} \in (0, 1)$ since then:

$$\begin{aligned} \sum_{t=1}^{\infty} \|\mathbf{z}^{(t)} - \mathbf{z}_{\text{exact}}^{(t)}\|_2 &\leq \sum_{t=1}^{\infty} \|\mathbf{v}^{(t)}\|_2 / \rho^{(t)} \\ &\leq \gamma_{\text{CG}} \sum_{t=1}^{\infty} \xi_{\text{CG}}^t = \frac{\gamma_{\text{CG}} \xi_{\text{CG}}}{1 - \xi_{\text{CG}}}. \end{aligned}$$

Note that the sum is finite if and only if $|\xi_{\text{CG}}| < 1$.

References

1. É. Thiébaud, “MiRA: an effective imaging algorithm for optical interferometry,” in “Astronomical Telescopes and Instrumentation,” , vol. 7013, F. D. Markus Schöller, William C. Danchi, ed. (SPIE, 2008), vol. 7013, pp. 70131I–1–70131I–12.
2. D. F. Buscher, “Direct maximum-entropy image reconstruction from the bispectrum,” in “IAU Symp. 158: Very High Angular Resolution Imaging,” , J. G. Robertson and W. J. Tango, eds. (1994), pp. 91–+.
3. F. Baron and J. S. Young, “Image reconstruction at Cambridge University,” in “Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series,” , vol. 7013 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference* (2008), vol. 7013 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, p. 144.
4. S. Meimon, L. M. Mugnier, and G. le Besnerais, “Reconstruction method for weak-phase optical interferometry,” *Optics Letters* **30**, 1809–1811 (2005).
5. K.-H. Hofmann and G. Weigelt, “Iterative image reconstruction from the bispectrum,” *Astron. Astrophys.* **278**, 328–339 (1993).
6. J.-B. le Bouquin, S. Lacour, S. Renard, E. Thiébaud, and A. Merand, “Pre-maximum spectro-imaging of the Mira star T Lep with AMBER/VLTI,” *Astron. Astrophys.* **496**, L1–L4 (2009).
7. F. Soulez, É. Thiébaud, S. Bongard, and R. Bacon, “Restoration of hyperspectral astronomical data from integral field spectrograph,” in “3rd IEEE WHISPERS,” (Lisbon, Portugal, 2011).
8. S. Bongard, F. Soulez, É. Thiébaud, and E. Pécontal, “3-D deconvolution of hyper-spectral astronomical data,” *Month. Not. Roy. Astron. Soc.* **418**, 258–270 (2011).
9. S. Bourguignon, D. Mary, and É. Slezak, “Restoration of astrophysical spectra with sparsity constraints: Models and algorithms,” *IEEE Journal of Selected Topics in Signal Processing* **5**, 1002–1013 (2011).

10. S. Gillessen, F. Eisenhauer, G. Perrin, W. Brandner, C. Straubmeier, K. Perraut, A. Amorim, M. Schöller, C. Araujo-Hauck, and H. Bartko, "Gravity: a four-telescope beam combiner instrument for the vlti," in "Proceedings of SPIE, the International Society for Optical Engineering," (Society of Photo-Optical Instrumentation Engineers, 2010).
11. J. A. Högbom, "Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines," *Astron. Astrophys. Suppl.* **15**, 417–426 (1974).
12. U. J. Schwarz, "Mathematical-statistical Description of the Iterative Beam Removing Technique (Method CLEAN)," *Astron. Astrophys.* **65**, 345–356 (1978).
13. F. Vincent, T. Paumard, G. Perrin, L. Mugnier, F. Eisenhauer, and S. Gillessen, "Performance of astrometric detection of a hotspot orbiting on the innermost stable circular orbit of the galactic centre black hole," *Monthly Notices of the Royal Astronomical Society* **412**, 2653–2664 (2011).
14. D. Donoho, "For most large underdetermined systems of linear equations, the minimal ℓ_1 norm near-solution approximates the sparsest near-solution," *Communications on Pure and Applied Mathematics* **59**, 907–934 (2006).
15. M. Fornasier and H. Rauhut, "Recovery algorithms for vector valued data with joint sparsity constraints," *SIAM Journal on Numerical Analysis* **46**, 577–613 (2008).
16. K. A. Marsh and J. M. Richardson, "The objective function implicit in the CLEAN algorithm," *Astron. Astrophys.* **182**, 174–178 (1987).
17. É. Thiébaud and F. Soulez, "Multi-wavelength imaging algorithm for optical interferometry," in "SPIE Conf. on Astronomical Telescopes and Instrumentation," , vol. 8445 (Amsterdam, 2012), vol. 8445.
18. É. Thiébaud, "Image reconstruction with optical interferometers," *New Astronomy Reviews* **53**, 312–328 (2009).
19. É. Thiébaud and J.-F. Giovannelli, "Image Reconstruction in Optical Interferometry," *IEEE Signal Process. Mag.* **27**, 97–109 (2010).
20. J. A. Fessler and B. P. Sutton, "Nonuniform Fast Fourier Transforms Using Min-Max Interpolation," *IEEE Trans. Signal Process.* **51**, 560–574 (2003).
21. M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**, 49–67 (2006).
22. M. Kowalski, "Sparse regression using mixed norms," *Applied and Computational Harmonic Analysis* **27**, 303 – 324 (2009).
23. R. Jenatton, J. Audibert, and F. Bach, "Structured variable selection with sparsity-inducing norms," *The Journal of Machine Learning Research* **12**, 2777–2824 (2011).
24. É. Thiébaud and L. Mugnier, "Maximum a posteriori planet detection and characterization with a nulling interferometer," in "IAU Colloq. 200: Direct Imaging of Exoplanets: Science & Techniques," , C. Aime and F. Vakili, eds. (2006), pp. 547–552.
25. J. Nocedal and S. J. Wright, *Numerical Optimization* (Springer Verlag, 2006), 2nd ed.
26. É. Thiébaud, "Optimization issues in blind deconvolution algorithms," in "Astronomical Data Analysis II," , vol. 4847, J.-L. Starck and F. D. Murtagh, eds. (SPIE, Bellingham, Washington, 2002), vol. 4847, pp. 174–183.
27. M. Ireland, J. Monnier, and N. Thureau, "Monte-Carlo Imaging for Optical Interferometry," in "Advances in Stellar Interferometry," , vol. 6268, J. D. Monnier, M. Schöller, and W. C. Danchi, eds. (SPIE, 2008), vol. 6268, pp. 62681T1–62681T8.
28. P. L. Combettes and J.-C. Pesquet, *Proximal splitting methods in signal processing* (Springer, New York, 2011), chap. Fixed-Point Algorithms for Inverse Problems in Science and Engineering, pp. 185–212.
29. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning* **3**, 1–122 (2010).
30. Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM Journal on Imaging Sciences* **1**, 248–272 (2008).
31. W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing," *SIAM Journal on Imaging Sciences* **1**, 143–168 (2008).
32. M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "An augmented lagrangian approach to the constrained optimization formulation of imaging inverse problems," *IEEE Trans. Image Process.* **20**, 681–695 (2011).
33. M. V. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Fast image recovery using variable splitting and constrained optimization," *IEEE Trans. Image Process.* **19**, 2345–2356 (2010).
34. M. R. Hestenes and E. Stiefel, "Methods of Conjugate Gradients for Solving Linear Systems," *Journal of Research of the National Bureau of Standards* **49**, 409–436 (1952).
35. J. Eckstein and D. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming* **55**, 293–318 (1992).
36. M. Figueiredo, R. Nowak, and S. Wright, "Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems," *IEEE Journal of Selected Topics in Signal Processing: Special Issue on Convex Optimization Methods for Signal Processing* **1**, 586–598 (2007).

37. S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. Signal Process.* **57**, 2479–2493 (2009).
38. G. H. Jacoby, D. A. Hunter, and C. A. Christian, "A library of stellar spectra," *Astrophys. J. Suppl.* **56**, 257–281 (1984).
39. M. Frigo and S. G. Johnson, "The Design and Implementation of FFTW3," *Proc. IEEE* **93**, 216–231 (2005). Special issue on "Program Generation, Optimization, and Platform Adaptation".
40. F. Delplancke, F. Deric, F. Paresce, A. Glindemann, F. Lévy, S. Lévêque, and S. Ménardi, "PRIMA for the VLTI - Science," *Astrophys. Space Sci.* **286**, 99–104 (2003).